European Journal of Operational Research 232 (2014) 342–349

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Determining near optimal base-stock levels in two-stage general inventory systems

Martin Albrecht *

Paul Hartmann AG, Paul-Hartmann-Str. 12, 89522 Heidenheim, Germany

ARTICLE INFO

Article history:
Received 9 April 2012
Accepted 17 July 2013
Available online 30 July 2013

Keywords:
Supply chain management
Inventory management
Safety stocks
General inventory system

ABSTRACT

This paper proposes an easily implementable, scalable decomposition heuristic for determining near optimal base stocks in two-level general inventory systems. In this heuristic, the general system is decomposed into assembly systems—one for each end product. For these assembly systems, the base-stock levels are calculated separately, taking into account risk-pooling effects for the common components. Our numerical analyses yield two main insights: First, the base-stock levels determined by the heuristic are close-to-optimal. Second, considerable improvements can be obtained compared to common-sense heuristics.

1. Introduction

This paper considers a two-level general inventory system. Outside suppliers deliver components to a raw material warehouse. Customers' demands are fulfilled from stocks of finished products. For replenishing these stocks, components are taken out of the warehouse and assembled. We have product-specific components as well as common components, which are included in the bill-of-materials of multiple end products. The assembly capacity is unlimited. The lead times for producing the finished product and for supply of components are deterministic and may be significant. There are no fixed costs for supply and assembly. Demands are stochastic. In case of shortages of finished products, demands are backordered. Both components and finished products can be held in stock at the expense of item-specific inventory holding costs. All problem parameters are stationary. The objective is to minimize the long-run average expected inventory and backorder costs.

This setting has a high practical relevance. Examples are the production of circuit boards (e.g., Grotzinger, Srinivasan, Akella, & Bollapragada, 1993) and of surgery sets (e.g., customer procedure trays at Paul Hartmann AG). Surgery sets consist of up to 60 different components. Customers can configure their sets on their own, thereby choosing between several thousands of different components. Often customers order their sets continuously during some years. Before a set can be sold, it has to be sterilized. Since sterilization and transport after assembly take at least 3 weeks (including quarantine time), stocks of finished products have to be built up. Capacities can often be regarded as unlimited in this context. Assembly personnel is extendable rather easily at short-term (e.g., through flexible working hours), and sterilization activities can be outsourced. As the (deterministic) sterilization and transport times are much larger than the assembly time, the assumption of deterministic production lead times seems appropriate in this context.

In practice, inventory management in such systems is often driven by simple heuristics without focus on cost-effectiveness. We believe that a major reason for this is the lack of scalable approaches which yield near optimal solutions and do not require heavy computation. Even the potential benefits from more sophisticated methods are largely unknown. Probably the most important topics in this context are the choice of the safety stocks and the strategy for component allocation. For component allocation, we assume a fixed and—in our opinion—favorable procedure. The main contribution of this paper is a new, simple heuristic for determining the safety stocks. We assume that the inventory levels of products and components are controlled using base-stock policies. Our heuristic decomposes the general system into separate assembly systems. For these assembly systems, the optimal base stocks can be determined using standard approaches. Among those is the procedure of Shang and Song (2003), which is the base for the heuristic developed by us. The main idea of our heuristic is to take the risk-pooling effects due to component commonality into account. When calculating the base stocks of the common components, the cumulated secondary demand for all their successors is considered. Apart from our new heuristic, we evaluate two "common-sense" heuristics in computational tests: level-by-level optimization and holding stocks of finished products only. Our heuristic performs well. In all of its variants, the resulting costs increased by less than 0.9–1.2% on the average, compared to the optimal solution from enumerating all combinations of base stocks

* Tel.: +49 7321 36 3058.
E-mail address: martin.albrecht@hartmann.info
URL: http://www.hartmann.info

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.07.025

between wide upper and lower limits. Compared to the common-sense heuristics, the improvement is considerable. The solutions have been more than 7.9% cheaper.

In the literature, effective heuristics for determining safety stocks in general inventory systems are rare. Despite the significant advances during the last years (e.g., Song & Yao, 2002; Benjaafar & ElHafsi, 2006; Song & Zhao, 2009; Lu & Song, 2005; Lu, Song, & Zhao, 2010; Dogru, Reiman, & Wang, 2010; Nadar, Akan, & Scheller-Wolf, 2012), this is also true for assemble-to-order (ATO) systems (see also Song & Zipkin, 2003 for this argument). ATO systems represent an important special case of the setting investigated in this paper: In ATO systems we have negligible lead times for assembly, such that finished products are not held in inventory. (A further related problem class is the repair kit problem (e.g., Smith, Chambers, & Shlifer, 1980). This problems share with ATO systems the assumption of zero finished product inventory since repair jobs cannot be done in advance. A main difference is that for component replenishment, the repairman undertakes a tour to a central depot. Thus, the repair kit models usually consider a single period or a periodic review policy with negligible component replenishment lead times.)

We are aware of two bodies of work dealing with scalable heuristics for general inventory systems (see also Graves & Willems, 2003 & de Kok & Fransoo, 2003 for overviews). The first is the so-called "guaranteed-service model approach" (e.g., Inderfurth & Minner, 1998; Graves, Willems, & Zipkin, 2000; Humair & Willems, 2006). Here a fixed external service time and upper bounded demand are assumed. Then each stage is able to quote a guaranteed, deterministic service time to its successors, which simplifies the determination of the safety stocks. The scope differs significantly from ours because the trade-off between inventory holding and backorder costs for finished products is not optimized by these models—this trade-off is exogenously predetermined by the amount of the external service time and the demand bound.

The second body corresponds to the "stochastic-service model approach". On the one hand, this includes large-scale, generic supply chain models where the stochastic lead times between stages are approximated by a random variable, while assuming that at most one component can be out of stock (e.g., Ettl, Feigin, Lin, & Yao, 2000). This assumption is especially critical in assembly structures with many components. There, we have a significant probability for stockouts of more than one component at one instance of time. Our heuristic does not rely on this assumption.

On the other hand, some researchers propose heuristics particularly designed for inventory systems with an assembly focus. Chew, Lee, and Lau (2010) determine safety stocks using infinitesimal perturbation analysis and a steepest descent algorithm. This requires significant problem-specific computing, in contrast to the heuristic proposed here. de Kok and Visschers (1999) propose a decomposition algorithm. For the special case where different components have the longest cumulative lead times, they outline a concrete solution procedure. More general settings, like the one investigated in this paper, are not considered. In this special case, the system is decomposed into serial systems, like our approach. The solution procedure of de Kok and Visschers (1999), however, cannot be compared with ours since its objective is different: the minimization of inventory subject to a service-level constraint. However, there is one important analogy: De Kok and Visschers solve the serial systems separately. Since they do not take risk-pooling effects into account, however, their solutions will often become suboptimal. We investigated an analogous scenario in our computational tests of Section 4 ($\alpha_i = 0$). There, the degree of suboptimality ranged between 1.8% and 5.1% compared to the preferred variant of our heuristic.

This paper is organized as follows. Section 2 presents the model and the allocation procedure. The decomposition heuristic is described in Section 3. The computational study in Section 4 provides insights into the solution quality of the heuristic and the potential savings compared to common-sense heuristics. Section 5 concludes the paper.

2. Model and allocation procedure

Consider a two-level general inventory system. There are $i = 1, \ldots, m$ finished products. The products are assembled on resources with unlimited capacities using components $j = 1, \ldots, n$. Denote the number of units of component j required to assemble product i by r_{ji}, where $r_{ji} \geqslant 0$. The components are procured from outside suppliers with unlimited capacities. The lead times L_i and L_j for assembly and supply are deterministic. Inventory holding is allowed for all items. Denote by h_i and h_j the (local) holding costs for the products and components, respectively, per unit time. The echelon holding costs of products are h_i^e. For each product i, customers' demand follows a Poisson process $\{D_i(t)\colon t > 0\}$ with rate λ_i, where $D_i(t)$ is the cumulative demand in the interval $(0,t]$. (Note that the heuristic is applicable for other distributions, too. Therefore, we will keep our exposition as generic as possible and highlight the settings specific to Poisson demand.) Demands are satisfied from on-hand inventory if possible. Shortages are backordered. Each unit backorder of product i incurs a (penalty) cost b_i.

The following variables are used to describe the dynamics of the system:

For the products,

$B_i(t)$	backorders of product i in period t,
$I_i(t)$	on-hand inventory of product i in period t,
$IN_i(t)$	net inventory of product i in period t, $IN_i(t) = I_i(t) - B_i(t)$,
$IO_i(t)$	inventory on order of product i in period t,
$IT_i(t)$	work-in-progress inventory of product i in period t (this corresponds to transit inventory in multiple-retailer models),
$ITP_i(t)$	work-in-progress inventory position of product i in period t, $ITP_i(t) = IT_i(t) + IN_i(t)$.

For the components,

$W_j(t)$	units of component j that are consumed for production in period t,
$I_j(t)$	on-hand inventory of component j in period t,
$IO_j(t)$	inventory on order of component j in period t.

Moreover, denote D_i as the lead time demand for product i. This is a random variable with the distribution of $D_i(t, t + L_i] = D_i(t + L_i) - D_i(t)$.

The sequence of events in the system is as follows: At first, the on-hand component inventories are allocated to the outstanding orders (a). The production of finished goods is carried out and inventory replenishments from the suppliers arrive (b). The customers' demands occur (c). The resulting inventories as well as new reorders are determined (d). Last, holding and backorder costs are calculated (e).

The allocation procedure, which determines $W_j(t)$ and, consequently, $IT_i(t)$ in step (a), will be described below. For actualizing the inventories in step (d), we rely on the following calculation. The net inventory of product i is determined considering the incoming and outgoing elements over the production lead time:

$$IN_i(t) = IN_i(t - L_i) + IT_i(t - L_i) - D_i \quad \text{for } i = 1, \ldots, m. \tag{1}$$